CHAPTER

6

REMOTE ACCESS IN WINDOWS 2000

After reading this chapter and completing the exercises, you will be able to:

- Describe the use of Routing and Remote Access Service (RRAS)
- Install RRAS
- Configure Inbound RRAS Connections
- Create a remote access policy
- Configure a remote access profile
- Configure a Virtual Private Network
- Configure remote access security, including encryption and authentication protocols
- Configure multilink connections
- Configure routing and remote access for DHCP integration
- Manage, monitor, and troubleshoot remote access

Remote access in Windows 2000 comes in the form of a service named the **Routing and Remote Access Service**, or RRAS. Essentially, this service runs on a Windows 2000 server and enables other servers or client computers that are not connected to the network via a permanent cable to establish temporary connections over phone lines, Integrated Services Digital Network (ISDN) lines, or services such as X.25, a standard protocol suite in a packet-switched network. Once a computer establishes a connection with the RRAS server, that computer can access the resources on the RRAS server and possibly access the other computers on the same network as the server, depending on the server's configuration.

This chapter begins with an overview of remote access that includes a brief history of remote access in the Windows environment and an examination of the many features and components that make up RRAS. You then learn how to install, configure, secure, and manage an RRAS server.

Remote Access Overview

Before you can understand the details of installing and configuring a remote access server, it is necessary to understand some of what goes on behind the scenes. Remote access uses many different protocols, including remote access protocols, networking protocols, and security protocols. This overview provides a look at the history, features, and concepts of remote access.

Brief History of Remote Access

The Remote Access Service, commonly called RAS, was first introduced with Windows NT 3.51 Service Pack 2 in an attempt to offer a simple and inexpensive way for remote users to dial in to a server and access network resources. This service was carried over almost fully intact to Windows NT 4.0.

Microsoft later introduced the Routing and Remote Access Service (RRAS), a substantial upgrade to RAS. Among other things, RRAS introduced the capability of multiprotocol routing to remote access. Previously, RAS supported only the NetBEUI networking protocol. For clients using other protocols, such as TCP/IP, RAS provided translation in the form of a NetBIOS gateway.

Windows 2000 includes significantly updated versions of RRAS. The Windows 2000 Server version now offers a number of new features:

- Internet Group Management Protocol (IGMP) support
- Network Address Translation (NAT), which allows computers on a LAN to share a single Internet connection
- Integrated AppleTalk routing
- Layer-Two Tunneling Protocol (L2TP) over IP Security (IPSec) support for router-to-router Virtual Private Networking (VPN) connections
- Improved support for Remote Authentication Dial-In User Support (RADIUS)

Routing and Remote Access Concepts

A Windows 2000 Server running the Routing and Remote Access Service can accept connections from users physically separated from the main network, but still needing to connect to the main network to access resources. Once connected, remote access clients use standard tools and applications to access these network resources. For example, once a user connects to a remote access server, that user can retrieve files with Windows Explorer, connect to a messaging server with a standard e-mail client, and open documents with applications like Microsoft Word. User's perceive themselves as directly connected to the network. In fact, at its heart, RRAS really is just another way to transmit standard networking protocols and commands already in use on the network. Instead of being put onto a network cable by a network interface card, the information is formatted (and possibly secured) by RRAS and transmitted across whatever type of link is configured.

Remote Access versus Remote Control

The concepts of remote access and remote control are often confused. Although both involve connecting to a remote computer, they are substantially different approaches.

In **remote access**, a client computer connects to a remote access server using a dial-up or other type of on-demand connection. Once connected to the network, the client can access network resources. All applications still run on the client computer.

In **remote control**, a client computer connects to a remote server and actually takes control over that server in a separate window on the client computer, as if the user were sitting at the server computer. All applications run on the server.

The Windows Routing and Remote Access Service does not support remote control. Remote control requires the use of Windows Terminal Service or third-party software like Symantec's pcAnywhere.

Remote Access Connection Types

Just like the RRAS service in Windows NT 4.0, Windows 2000 RRAS provides two distinct types of remote access connections to remote users:

- Dial-Up Networking
- Virtual Private Networking

Dial-Up Networking With **Dial-Up Networking**, a client makes a temporary, dial-up connection to a physical port on the RRAS server. This connection uses the services of a public telecommunications provider, such as an analog phone line, an ISDN line, or X.25. A good example of dial-up networking is when both client and server have a standard modem. The client initiates the dial-up connection using the modem and makes the connection to the server modem over public phone lines. The server authenticates the user and provides the configured access.

Virtual Private Networking **Virtual Private Networking (VPN)** provides a way of making a secured, private connection from the client to the server over a public network such as the Internet. Unlike dial-up networking, where client and server share a direct physical connection, a VPN connection is logical and not necessarily direct. Typically, a remote user connects to an Internet Service Provider (ISP) using a form of dial-up networking, though establishing a VPN connection over a standard cable-connected network is quite possible. The RRAS Server also connects to the Internet (probably via a persistent, or permanent, connection) and is configured to accept VPN connections. Once connected to the Internet, the client then establishes a VPN connection over that dial-up connection to the RRAS server.

VPN offers two significant advantages. First, remote users who are not in the same local calling area as the remote access server need not make long distance calls to connect to the network. Instead, they can make local calls to an ISP. Second, every standard dial-up

connection requires that a physical device be present on the RRAS server and devoted to that connection. This limits the number of users that can connect remotely at a single time. Assuming a fairly high-bandwidth Internet connection from the RRAS server to the Internet, more remote users can connect at the same time using VPN than users with dial-up connections.

Protocols

A **networking protocol** is simply a defined and often standardized way of communicating between two devices on a network. You must be familiar with two general types of protocols to work with RRAS: remote access (or line) protocols and networking (or LAN) protocols.

Remote Access Protocols **Remote access protocols** govern how information is broken up and transmitted over wide area network (WAN) connections, of which a dial-up connection is one type. RRAS supports four remote access protocols:

- **Point-to-Point Protocol (PPP)**: an industry standard set of robust and flexible protocols, by far the most common remote access protocol used today. Most dial-in servers, including RRAS, support PPP, and it is generally considered the best choice for remote access situations. Windows 2000 RRAS supports PPP both for dial-out and dial-in connections.
- **Serial Line Interface Protocol (SLIP)**: an older protocol developed in UNIX and still widely used. Windows 2000 RRAS supports SLIP in dial-out configurations, but you cannot use a SLIP client to dial in to an RRAS server.
- RAS Protocol: a proprietary protocol, used only between Microsoft-based networks, that supports the NetBIOS naming convention. It is required to support NetBIOS naming and is installed by default when you install the RRAS server.
- NetBIOS Gateway: provides compatibility with older versions of RAS Server that do not support networking protocols such as TCP/IP and NWLink. The NetBIOS gateway translates data from the NetBEUI protocol to these other protocols.

Networking Protocols Networking protocols govern how information is transmitted between devices on a local area network (LAN). You can find detailed information on supporting networking protocols in Windows 2000 in Chapter 2, but a brief recap is in order here.

RRAS supports the use of three networking protocols:

- ***NetBEUI***: a simple, efficient protocol primarily used on small networks that consist only of Microsoft clients. Although easy to configure and manage, NetBEUI does not support routing and is therefore not suitable for large, varied networks.
- ***Transmission Control Protocol/Internet Protocol (TCP/IP)***: an extensive, robust protocol ideally suited for connecting different types of computers and operating systems. Thus, it is the standard choice of protocols for networks containing many different types of systems, such as Microsoft systems or those based on UNIX, and it is the standard protocol for the Internet.

- ***Internetwork Packet eXchange (IPX)***: the protocol of choice for networks using Novell's NetWare. If your network uses NetWare and your remote clients need to access these resources, you must enable IPX.

You must choose at least one LAN protocol to use, but RRAS enables you to use all three simultaneously if necessary. Keep in mind that any remote client dialing in must support one of these protocols. RRAS also supports the Point-to-Point Tunneling Protocol (PPTP), an extension of PPP that you can use to establish a connection such as a Virtual Private Network.

Remote Access Clients

Just about any dial-up client software that supports PPP can connect to an RRAS server. Such clients include:

- Windows 2000
- Windows NT 4.0
- Windows NT 3.5
- Windows 95/98/ME
- Windows for Workgroups 3.1x
- MS-DOS
- Microsoft LAN Manager remote access clients
- UNIX and Apple Macintosh clients using third-party client software

Remote Access Features

Now that you are familiar with some of the concepts involved in remote access, it's time to look at some of the other features offered by RRAS in Windows 2000. The updated version provides new support for router discovery, network address translation, multicast routing, and powerful remote access policies.

Router Discovery

Windows 2000 supports a new feature called router discovery that provides a method for detecting default gateways. Manual gateway configuration or DHCP offers clients no way to adapt to changes in network configuration. Using router discovery, however, clients can dynamically determine the status of routers and switch to back-up routers, should a primary router fail. Chapter 7 covers this feature in more detail.

Network Address Translator

Network Address Translator (NAT) is a router standard that translates IP addresses on a private network into valid Internet IP addresses. This means that a single computer with Internet connectivity can share its Internet connection with other computers on the network

through a single IP address. Windows 2000 Server features both a full-featured NAT named Connection Sharing and an easier to configure version named Shared Access. You learn about both later in this chapter.

Multicast Routing

Multicast routing, or multicasting, is a targeted form of network broadcasting that sends information to a select group of users instead of all users connected to a network. Standards are being developed to support multicasting over a TCP/IP network such as the Internet.

Windows 2000 Server supports multicast routing using what is known as a multicast proxy. You can use this proxy to extend multicast support to remote access users or to a single LAN network connected to the Internet. Windows 2000 Server behaves like a multicast client, communicating between local multicast clients on the LAN and multicast servers on the Internet.

Remote Access Policies

In earlier versions of Windows NT, remote access was granted to users based on a single option configurable in User Manager or the Remote Access Admin tool. With this option, named Grant Dial-In Permission To User, enabled for a user account, that user could dial in to the remote access server.

In Windows 2000, granting remote access privileges is more flexible and more complex. Each User object in the Users and Computers tool (or the Active Directory Users and Computers tool if a member of a domain) has certain dial-in properties.

In addition to the dial-in properties of a user, Remote Access Policies (RAPs) are used to configure conditions under which a user may connect using a specific remote access connection. Administrators may now include restrictions based on criteria such as time of day, type of connection, authentication, and even length of connection. You learn more about configuring RAPs later in this chapter.

Remote Access Security

Remote access has always been considered one of the weaker points of networking security. While it's always been fairly easy to secure a network from unauthorized physical access, the current popularity of Internet access and remote user access places larger security demands on the modern network. Fortunately, new security technologies and protocols have been developed to ease the problem of remote access security.

User Authentication

The primary method of securing a remote access connection involves authenticating the user trying to connect. To do this, the user (or the user's client computer) must present some sort of credentials that allow the RRAS server to verify that the user is indeed a valid user.

Windows 2000 supports five different user **authentication** protocols:

- Password Authentication Protocol (PAP)
- Shiva Password Authentication Protocol (SPAP)
- Challenge Handshake Authentication Protocol (CHAP)
- Microsoft CHAP (MS-CHAP)
- Extensible Authentication Protocol (EAP)

Password Authentication Protocol Password Authentication Protocol is the most basic form of user authentication. A user's name and password are transmitted over the dial-up connection to the RRAS server. Transmitted in clear text with no encryption, this information is quite vulnerable to snooping. In addition, PAP provides no way for a client and server to authenticate one another. For the most part, better authentication protocols have rendered PAP obsolete. In fact, Microsoft recommends that you do not use it unless absolutely necessary.

Shiva Password Authentication Protocol Windows 2000 Server includes Shiva Password Authentication Protocol mainly for compatibility with remote access hardware devices manufactured by Shiva, a private company now owned by Intel. SPAP isn't used much on most networks.

Challenge Handshake Authentication Protocol Considerably more secure than PAP or SPAP, Challenge Handshake Authentication is a form of authentication in which the server sends the client a key to encrypt the client's username and password. The client then sends the encrypted information across the dial-up connection to the server, which decrypts it and attempts to validate the user. The encrypted username and password are considerably less vulnerable to eavesdroppers. CHAP is also commonly called MD5-CHAP because it uses the RSA MD5 hash algorithm for encryption.

Microsoft CHAP A modified version of CHAP, Microsoft CHAP allows the use of Windows 2000 authentication information. Of the two versions of MS-CHAP, version 2 is most secure; all Microsoft operating systems support it. Other operating systems sometimes support version 1.

Extensible Authentication Protocol A general protocol for PPP authentication, Extensible Authentication Protocol supports multiple authentication mechanisms. Instead of selecting a single authentication method for a connection, EAP can negotiate an authentication method at connect time. The computer asking for the authentication method is called the authenticator and may require several different pieces of authentication information. This allows the use of almost any authentication method, including secure access tokens or one-time password systems.

Each authentication method supported in EAP is called an EAP type. Both the client and the server must support the same EAP type. Currently, Windows 2000 comes with the following two EAP types:

- *EAP MD5-CHAP*: virtually identical to the normal CHAP authentication, except that it packages and sends authentication information as EAP messages. This means that if you turn on EAP MD5-CHAP and disable regular CHAP on the server, regular CHAP clients cannot connect.
- *EAP Transport Level Security (TLS)*: lets you use public-key certificates for authentication. TLS is similar to Secure Sockets Layer (SSL) used in most Web browsers. With EAP TLS, both the client and the server send encrypted authentication messages. EAP TLS is one of the strongest authentication methods available but requires that your RRAS server belong to a Windows 2000 domain.

A third EAP authentication method is included with Windows 2000, although it is not technically an EAP type. EAP-RADIUS passes authentication information to a RADIUS server for authentication. A RADIUS server is usually devoted to running a large user account database against which it can identify remote users. In addition, RADIUS can authenticate users from a wide variety of accounts, including Windows domains, Novell Directory Services, SQL Server databases, and UNIX password files.

Connection Control

In addition to its ability to authenticate users in a variety of ways, Windows 2000 RRAS provides a number of methods for securing the actual connection from a client to a server. One such method, Callback Control Protocol, allows your RRAS servers or clients to negotiate a callback with the other end. For example, you may configure a server to hang up and call a user back at a specified number whenever that user tries to connect. This provides two advantages. First, a successful connection can only be made from a particular number—a good way of ensuring that only authorized users can make the connection. Second, for users dialing in from another calling area, the company can foot the bill for the long-distance call.

Another way to control connections is to configure an RRAS server to accept or reject calls based on Caller ID or Automatic Number Identification (ANI) information. For example, you could configure a server to accept calls only from a certain number or to reject calls from callers who may be trying to break into the system.

Access Control

RRAS supports a number of ways to control remote user access to the RRAS server. The primary access control method is enabling or disabling permission to dial in on individual user accounts. In addition to this basic method, RAPs allow you to extend control over whether users can dial in or not by setting a number of conditions on the access.

Installing and Configuring Routing and Remote Access

Now that you have a basic grasp of the concepts behind the Routing and Remote Access Service, it is time to see how to actually implement and configure it. This is a fairly simple procedure, but unlike other software or Windows components, you cannot install RRAS using the Add/Remove Programs component of the Windows Control Panel. Instead, RRAS is automatically installed along with Windows 2000 Server but left in a disabled state. All you need to do is enable it. First, however, you must make sure that all dial-up equipment, interfaces, and protocols that you intend to use with the server are installed and configured correctly.

This section provides an overview of the set-up process and the choices you must make. Hands-on Project 6-1 at the end of this chapter walks you through the actual steps of setting up RRAS.

First, you log on to the server with Administrator privileges and open the Routing and Remote Access utility from the Administrative Tools program group on the Start menu. This utility (shown in Figure 6-1) is actually a snap-in for the Microsoft Management Console that controls most management features of Windows 2000.

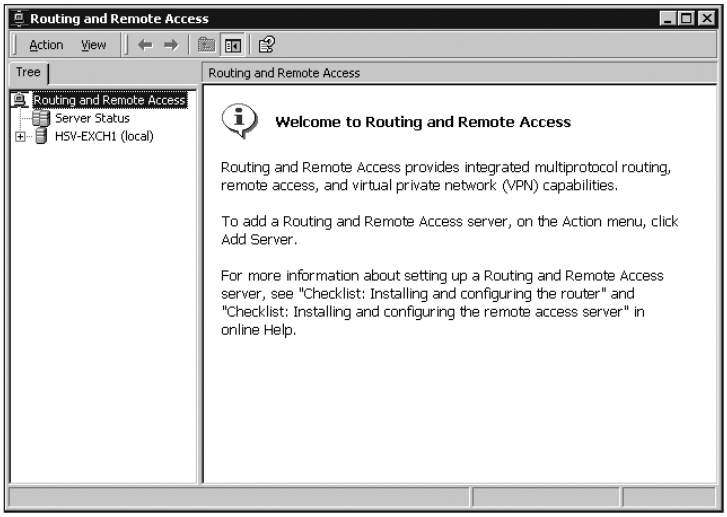

Figure 6-1 Routing and Remote Access snap-in

In the tree in the left pane, find and right-click the name of the server. From the shortcut menu that appears, choose the Configure and Enable Routing and Remote Access command to begin the Routing and Remote Access Server Setup Wizard. The wizard takes you through several configuration steps. The first, shown in Figure 6-2, asks you to select the type of configuration you want to install.

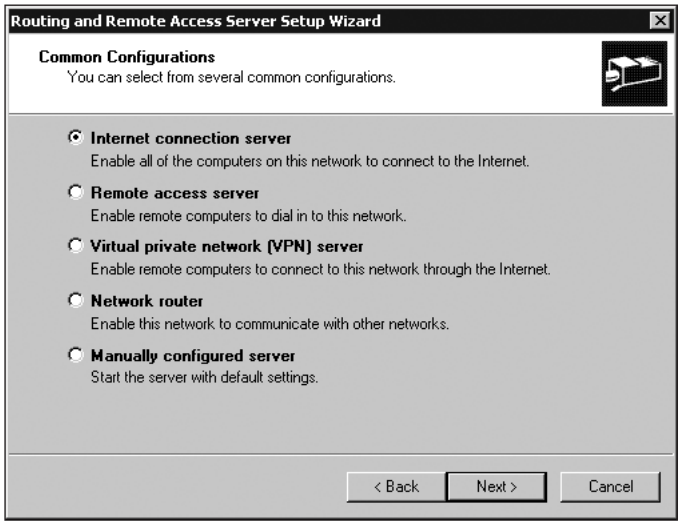

Figure 6-2 Routing and Remote Access Server Setup Wizard

Next, the RRAS Setup Wizard asks you to verify that the protocols you wish to use on the server are already installed and configured. If not, you must configure them before taking the next step.

Next, you configure some options for your network. These options include:

- Selecting the network adapter that you want to use on your internal network, as shown in Figure 6-3. This step is particularly important for a multi-homed server.

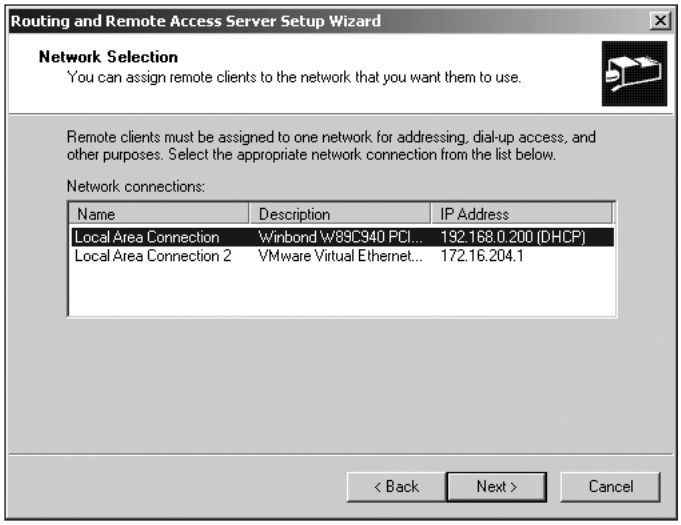

Figure 6-3 Selecting a network adapter

- Deciding whether to use DHCP or to define a static pool of IP addresses, as shown in Figure 6-4. If you have a DHCP server on the same subnet as the RRAS server, the DHCP option is usually best. You learn more about using RRAS and DHCP later in this chapter. If you choose to use a static pool of addresses, you are asked to configure the starting and ending IP addresses of the range you want to use.

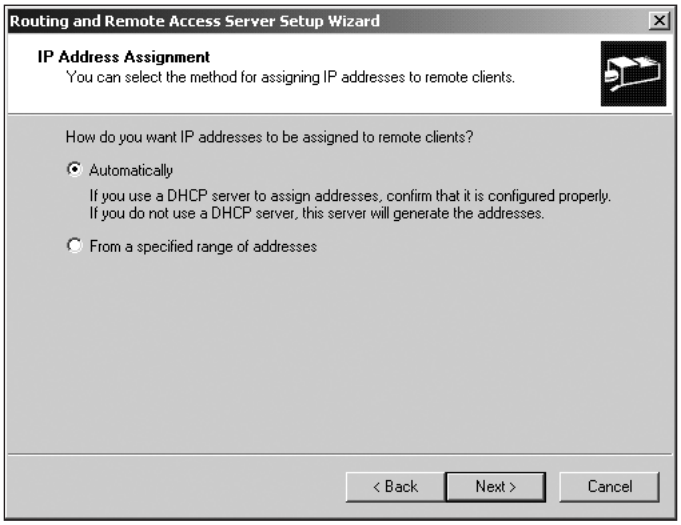

Figure 6-4 Configuring IP addressing for remote clients

- Deciding whether to use Windows authentication or RADIUS. Choosing RADIUS configures RRAS to be a RADIUS client using the EAP authentication protocol, as discussed previously in this chapter. You are also asked to configure settings for the RADIUS server.

And that's all there is to it. Once you enable RRAS, you can pause or stop the service by right-clicking the server in the Routing and Remote Access snap-in and choosing the appropriate action from the All Tasks menu on the shortcut menu.

The following section explains how to configure connections for the new server.

Configuring Remote Access

Once you enable RRAS on your server, you need to configure it to behave the way you want. This configuration occurs in three places:

- Most configuration of inbound connections happens at the server level using the Routing and Remote Access Service snap-in that you used to enable the service. In particular, you use the server object's property page to control

whether the server allows connections at all, what protocols it supports and how, security options, and event logging. You also use RRAS to set policies and profiles and to monitor the status of a remote access server.

- A good bit of configuration also happens using the property pages for individual users in the Active Directory Users and Computers snap-in. Here you grant dial-in permissions for individual users, as well as set callback and other dial-in options.
- Once you configure the server and the user accounts for dial-in access, you also need to configure each client. Fortunately, all versions of Windows and most other operating systems come with some built-in form of dial-in capability that is relatively easy to configure. In Windows, this capability is named dial-up networking.

The following section discusses the first two types of configuration. Configuring clients is very similar in different versions of Windows, so we do not go into detail here. Hands-on Project 6-6 at the end of the chapter provides some practice configuring a dial-up connection in Windows 2000 Professional.

Configuring Inbound Connections on the Server

Like most objects in Windows 2000 management, RRAS servers are configured using the settings on a number of property pages. You can access these pages by right-clicking on the server and choosing the Properties command from the shortcut menu. This section explains the general use of each of these property pages.

General Properties

The first page you see when you open the properties for an RRAS Server is the General page, shown in Figure 6-5. The most important setting on this page is the **Remote access server** option, which allows the RRAS Service to operate as a remote access server. This means you can switch remote access on and off without actually stopping the RRAS service, an action that causes the service to erase its settings.

The other option on the General page, **Router**, lets you choose whether clients accessing the RRAS Server can also access the rest of the network that is connected to the server. With routing enabled, you can choose whether to allow routing access only to the LAN (computers that are directly connected to the computer) or to allow demand-dial routing as well. **Demand-dial routing** allows the RRAS server to make WAN connections to other remote networks.

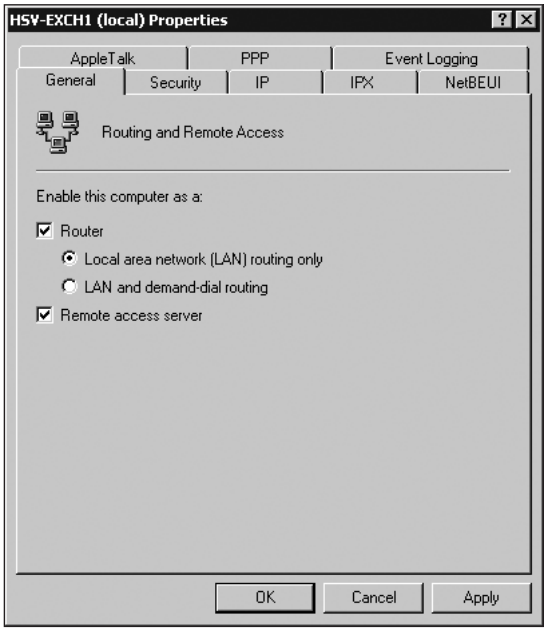

Figure 6-5 General property page for an RRAS server

Security Properties

You use the Security page, shown in Figure 6-6, to specify which authentication and accounting methods RRAS uses. You can choose one of two authentication providers from the list on this page:

- *Windows Authentication*: indicates the built-in authentication suite provided with Windows 2000. This suite includes several authentication protocols, including PAP, SPAP, CHAP, MS-CHAP, MS-CHAP v2, and EAP, all of which were discussed in detail earlier in the chapter. If you choose the Windows Authentication provider, you can click the Authentication Methods button to open a separate dialog box that lets you enable or disable individual authentication protocols.
- *RADIUS Authentication*: causes all authentication requests to be forwarded to a RADIUS server for approval. This authentication method is also discussed earlier in the chapter. If you choose the RADIUS Authentication method, you can click the Configure button to the right of the drop-down list to set up communications with the RADIUS server.

The **Accounting provider** list on the Security page allows you to configure whether connection request events are sent to the Windows Event Log (the Windows Accounting option) or to a RADIUS server (the RADIUS accounting option).

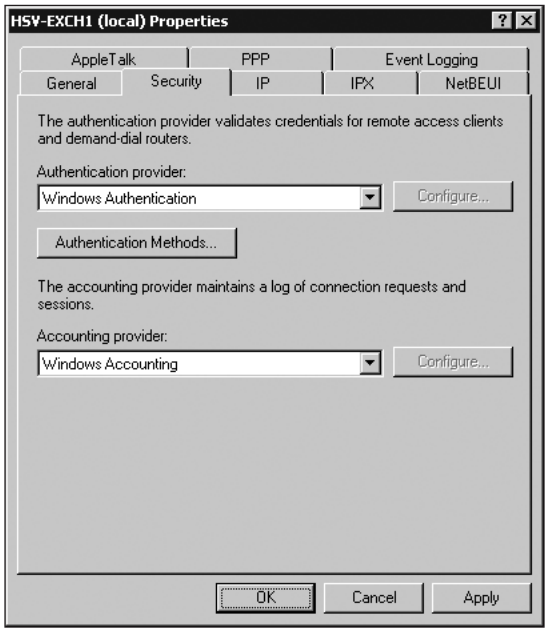

Figure 6-6 Security property page for an RRAS server

PPP Properties

Figure 6-7 shows the PPP property page used to control the PPP-layer options available to clients.

Options on this page include:

- *PPP* ***Multilink Protocol (MP)*** combines multiple physical links into a single logical link. For example, you could combine two 56-KB modem links into a 128-KB link. The multilink protocol is turned on by default, but if your clients are not using it (or your server does not support multiple physical connections), there is really no reason to leave it turned on. Hands-on Project 6-2 at the end of the chapter details steps for disabling the Multilink protocol.
- ***Bandwidth Allocation Protocol (BAP)*** *and* ***Bandwidth Allocation Control Protocol (BACP)***: allow a client to add and remove links dynamically during a multilink session to adjust for changes in bandwidth needs. This option is available only when you enable the Multilink option.

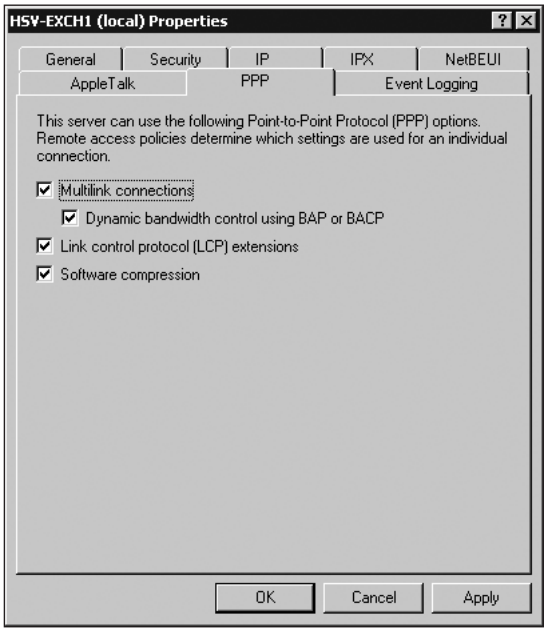

Figure 6-7 PPP property page for an RRAS server

- **Link control protocol (LCP)** extensions include a number of enhancements to the LCP protocol that establishes a PPP link and controls its settings. One of the primary enhancements included is the ability of the client and server to agree dynamically on protocols used on the connection. This option is turned on by default and, since Windows 9x, NT, and 2000 clients all support the extensions, you probably want to leave it that way.
- The Software compression option controls whether RRAS should allow clients to use the Compression Control Protocol (CCP) to compress PPP traffic. Again, this option is on by default, and leaving it that way is usually best.

IP Properties

The IP property page is one of four property pages that control the networking protocols supported by RRAS. Others include IPX, NetBEUI, and AppleTalk. These pages only show if the protocols were installed on the server before enabling RRAS. To install protocols after RRAS, you must disable and re-enable RRAS. Figure 6-8 shows the IP property page, which controls the following properties:

- *Enable IP Routing*: controls whether RRAS routes IP packets between the client and the rest of the network that the RRAS server is on. Enabled by default, this option gives all TCP/IP-based clients access to the network. When disabled, clients can only access resources located on the RRAS server itself. This option depends on the Router option that you enable on the General property page for

the server. Since you can turn on routing for the server itself and turn off routing for the IP and other protocols individually, you can achieve pretty fine control over what clients can and cannot access on the network.

- *Allow IP-based remote access and demand-dial connections*: controls whether TCP/IP-based clients can connect to the RRAS server at all. Disabling this option makes the rest of the settings on this page moot (which may be a good choice if you want to allow access only to clients using other protocols).
- *IP address assignment*: controls how remote clients get their IP addresses when connecting to the RRAS server. The default setting is based on the answers you give when you work with the RRAS Setup Wizard. DHCP allows your RRAS server to refer clients to a DHCP server on your network to be assigned IP addresses dynamically. The section, "Configuring RAS for DHCP Integration," later in this chapter provides more information on DHCP. The Static address pool option lets you manually configure a pool of IP addresses that the RRAS server itself can assign to clients.

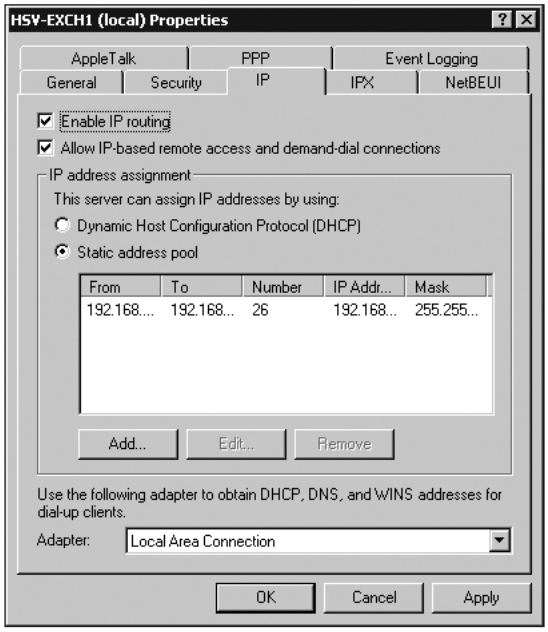

Figure 6-8 IP property page for an RRAS server

To practice configuring the protocols on an RRAS server, be sure to complete Hands-on Project 6-3 at the end of the chapter.

IPX Properties

IP and IPX are fairly similar protocols, even though some of their configuration details differ. On the IPX property page, shown in Figure 6-9, you can choose the following settings:

- *Allow IPX-based remote access and demand-dial connections*: specifies whether IPX-based clients can access the RRAS server at all.
- *Enable network access for remote clients and demand-dial connections*: works like the Enable IP Routing option on the IP page—it allows clients to access the network to which the RRAS server connects instead of just the RRAS server itself.
- *IPX Network Number Assignment*: controls the assignment of IPX addresses to remote clients. With the default option, the server takes care of this task automatically. This choice is probably the wisest, unless you have a specific reason for needing to assign addresses manually. You should also use the default setting, use the same network number for all IPX clients, so that clients can use all IPX resources. The Allow remote clients to request IPX node number option essentially lets remote clients configure their own addresses. Disabled by default, this option is often considered a security-risk: a client could misrepresent itself as another client.

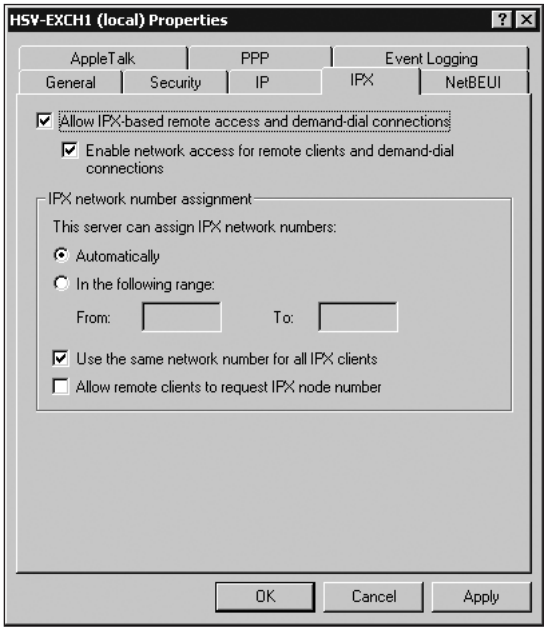

Figure 6-9 IPX property page for an RRAS server

Notice that both the IP and IPX property pages have an option that allows the client to access the rest of the network. However, also notice that the labels for these options are quite different for each protocol. You should be familiar with the wording of these labels for the exam.

NetBEUI and AppleTalk Pages

Both the NetBEUI and AppleTalk protocols are pretty simple, as exemplified by their property pages in RRAS. The NetBEUI protocol has an option for enabling the protocol and an option for whether clients can access only the server or the rest of the network as well. The AppleTalk page has only a setting for enabling the protocol.

Event Logging Page

The Event Logging property page lets you control the level at which events are logged either to the Windows Event Log or to a RADIUS server. The section, "Managing, Monitoring, and Troubleshooting RAS," discusses this page.

Configuring a User for Remote Access

Once you configure your RRAS server to support remote access, the next step is to configure what users can use that access and how. You will use three tools for doing this:

- ***User profiles***: configuration settings associated with individual user accounts. Each user has exactly one profile, usually stored in the Active Directory. These profiles include options such as whether the user can connect remotely, whether callback for the user is enabled, and so on.
- ***Remote access policies***: connection rules that apply to groups of users.
- ***Remote access profiles***: associated with policies and containing settings that determine what happens during call setup and completion.

Configuring User Profiles

If your RRAS server is part of a Windows 2000 domain, the Active Directory, a central directory of resources and objects for the entire network, stores users' profiles. In this case, you use the Active Directory Users and Computers snap-in to manage these profiles. If your RRAS server is not part of a Windows 2000 domain, the local computer stores users' profiles and the Local Users and Groups snap-in controls them. Whichever tool you use, the configuration of the user profiles is the same. In this chapter, we assume that you are using Active Directory Users and Computers.

Each user profile has a host of settings scattered across a number of property pages. However, in relation to remote access, you should be most concerned with the Dial-in page shown in Figure 6-10.

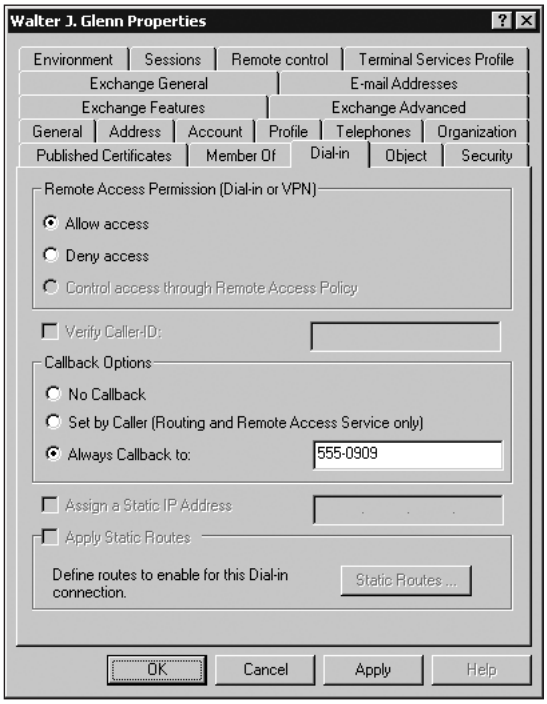

Figure 6-10 Dial-in property page of a user profile

A number of settings on this page control the user's remote access capability:

- *Remote Access Permission (Dial-in or VPN)*: enables remote access for the user. There are three options. The first two explicitly allow or deny access; these options override any policy settings that may apply to the user. The final option is to let a remote access policy control the user's access.
- *Verify Caller-ID Option*: lets you enter a phone number that is used to verify the remote caller using the Caller-ID information provided by the phone company. Calls from any other number are automatically rejected.
- *Callback Options*: provides two ways to use the callback feature of RRAS that lets the server automatically hang up on an incoming call and call the user back. The first option lets the caller set the number in the client software. This option does not add much security but can be a good way to let a company be billed for long-distance access instead of the caller. The second callback option sets a specific number that is always called for the user. This method adds some security, in that the user must call from a given number in order to access the network.
- *Assign a Static IP Address*: provides a user with the same IP address every time the user calls in. While it is generally a better idea to let RRAS work in conjunction with DHCP to assign dynamic IP addresses (covered later in the chapter), occasionally a client may need a static IP address for specific applications.

- *Apply Static Routes*: lets you define a set of routes always used to deliver information from the client to specified hosts on the network. If you do not enable this option, the client uses the default gateway assigned by DHCP or given manually. Once you enable this option, use the Static Routes button to add and remove routes.

Configuring Remote Access Policies

While user profiles define settings for an individual user, remote access policies define settings for a whole group of users. A policy is a set of rules that the system evaluates when it determines whether a user can access the network or not. User profiles and policies work together to provide dial-in capability. You can use a policy to define overall settings for a group of users, but individual settings in a user's profile override any policies in effect when that user logs on.

You manage remote access policies with the Routing and Remote Access Service snap-in through a container named Remote Access Policies, shown in Figure 6-11. As you can see, the only policy listed in the container by default is named Allow access if dial-in permission is enabled. This most basic of policies simply tells the RRAS service that if a user's profile grants dial-in access, it may grant that user remote access to the server.

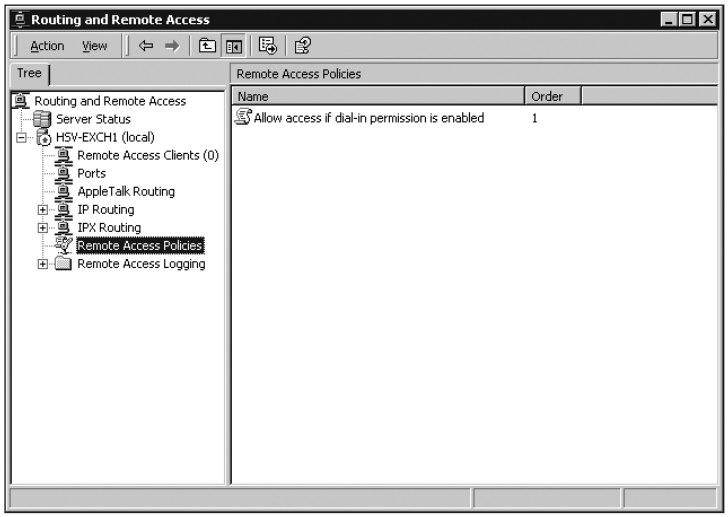

Figure 6-11 Remote Access Policies container

Creating a New Policy

To create a new policy using the RRAS snap-in, right-click the Remote Access Policies container and select the New Remote Access Policy command from the shortcut menu. This launches the Add Remote Access Policy Wizard.

You can practice using this wizard to create a policy in Hands-on Project 6-4. This section describes the general configuration of a new policy.

The first step the wizard takes you to is naming the policy. Once you do this, you see a page that lists the conditions for the new policy. Initially, this page is blank, but each new condition you add updates the list, as shown in Figure 6-12.

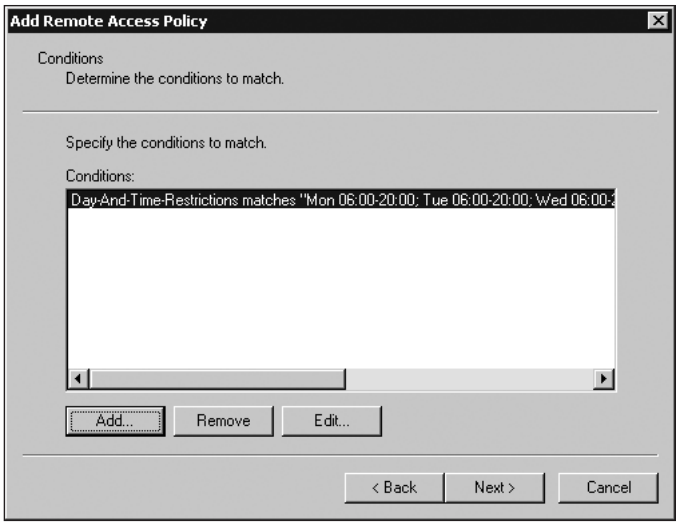

Figure 6-12 Setting conditions for a policy

To add a new condition, click the Add button. This opens the Select Attribute dialog box, which lists all available conditions. Table 6-1 details these conditions. Once you pick a condition from the list, a dialog box opens that lets you set configuration parameters that vary depending on the type of condition you choose. For example, choosing the Day-And-Time Restriction opens a dialog box that lets you restrict access by picking dates and times from a calendar.

Table 6-1 Remote access policy conditions

Attribute Name	What It Specifies
Called-Station-ID	Phone number dialed by user
Calling-Station-ID	Caller's phone number
Client-Friendly-Name	Name of the RADIUS server attempting to validate connection (IAS only)
Client-IP-Address	IP address of the RADIUS server attempting to validate connection (IAS only)
Client-Vendor	Manufacturer of RADIUS proxy or NAS (IAS only)

Table 6-1 Remote access policy conditions (continued)

Attribute Name	What It Specifies
Day-and-Time Restriction	Time periods and days during which connection attempts are accepted or rejected
Framed-Protocol	Remote access protocol (PPP, SLIP, and so on) used for framing incoming packets
NAS-Identifier	Name of the NAS that accepted the original connection (IAS only)
NAS-IP-Address	IP address of the NAS that accepted the original connection (IAS only)
NAS-Port-Type	Physical connection type (phone, ISDN, and so on) used by the caller
Service-Type	Type of service the user requested; types include framed for PPP or login for telnet
Tunnel-Type	Tunneling protocol that should be used (L2TP or PPTP)
Windows-Groups	Windows groups to which the user belongs

After choosing the conditions you want to apply in the new policy, the next step in creating a policy is to choose whether the policy is to allow users to connect or deny them connection. Each policy you create serves only one purpose.

The final step in the Add Remote Access Policy Wizard allows you to modify the remote access profile attached to the policy. You can do this when you establish the policy or come back to it later if you want. Configuring remote access profiles is discussed a bit later in the chapter. Once you finish the wizard, you have created the new remote access policy, you return to the RRAS snap-in, and the policy goes into effect.

Configuring Existing Policies Policies are evaluated in the order that they appear in the Remote Access Policies container in the RRAS snap-in. You can rearrange this order by right-clicking any policy and using the Move Up and Move Down commands on the shortcut menu. Order is very important, as each condition of each policy is considered to determine whether a user can access the system. All conditions of all policies must be met before access is granted.

Aside from ordering the policies, you can open the properties for any particular policy by right-clicking it and selecting Properties from the shortcut menu. The policy object has only one property page, which is shown in Figure 6-13. On this page, you can change the name of the policy, add new conditions to the policy, switch between granting and denying access based on those conditions, and edit the remote access profile for a policy.

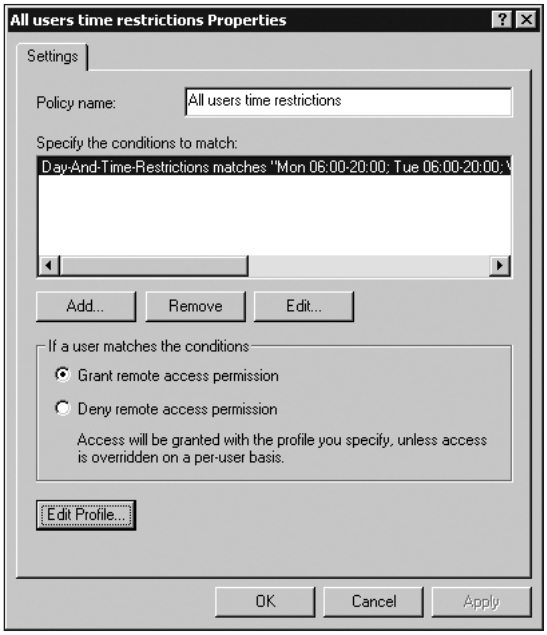

Figure 6-13 Editing properties for a condition

Configuring Remote Access Profiles

Remote access profiles are an important part of a good remote access policy strategy. The first thing, of course, is not to confuse remote access profiles with user profiles. User profiles, which we covered previously, are the collections of settings that pertain to an individual user and are stored in the Active Directory. Remote access profiles determine the remote access settings that apply to users when they meet the conditions in a policy and receive access. Each policy has one associated profile. You can open and edit the profile for a policy on the last page of the Add Remote Access Policy Wizard or later by using the property page for the policy. Either way, you click the Edit Profile button to begin making changes.

The remote access profile, sometimes referred to as a dial-in profile, has six tabs: Dial-In Constraints, IP, Multilink, Authentication, Encryption, and Advanced. The following sections explain these pages.

Dial-in Constraints Properties The Dial-in Constraints page, shown in Figure 6-14, sports a number of general dial-in controls. These controls let you drop a user if a connection remains idle for a certain time, restrict the maximum session length, restrict access to specified days and times, restrict access to a particular number, and even restrict the dial-in media types (ADSL, ISDN, and so on) allowed.

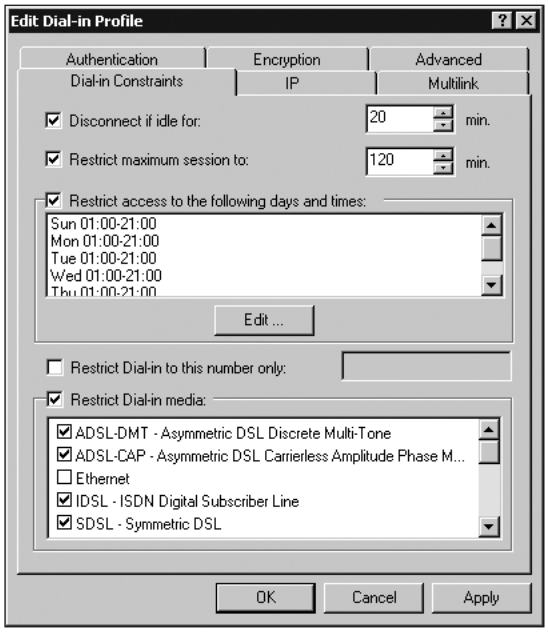

Figure 6-14 Dial-in Constraints page for a profile

IP Properties The IP page, shown in Figure 6-15, lets you control the IP settings for incoming connections. You can make two settings on this page. The IP Address Assignment Policy section controls how the client is assigned an IP address when connecting. Remember that these settings apply to all users granted access based on the policy to which the profile is attached. Configurations made to individual users (such as a static IP address) override these settings. The IP Packet Filters section lets you add advanced filters to prevent the client or the server from sending certain types of IP packets.

Multilink Properties Options on the Multilink page control how a client can connect using the Multilink Protocol and the Bandwidth Allocation Protocol. You can explicitly disable or allow multilink or you can set it to follow the default settings used for the server. The BAP settings allow you to specify the idle bandwidth threshold at which the number of lines in use is reduced.

Although the multilink properties suggest otherwise, any settings that you assign using a profile are not used at all unless the server settings match. For example, if you want to set the profile to allow multilink, then both multilink and BAP must be turned on at the server. Otherwise, the settings are ignored. This is true of other profile settings such as IP and authentication, as well.

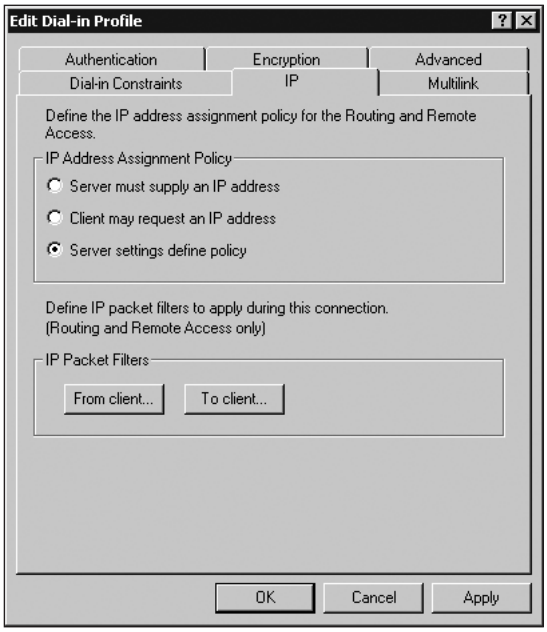

Figure 6-15 IP page for a profile

Authentication Properties The Authentication page lets you specify the authentication methods used for the policy attached to the profile. Simply select any authentications that you want the profile to allow. Remember that any methods you select must also be enabled at the server.

Encryption Properties The Encryption page, shown in Figure 6-16, lets you enable certain types of encryption for use on the connection. For some reason, though, Microsoft labeled these encryption methods as No Encryption, Basic, Strong, and Strongest instead of indicating the actual encryption algorithms used. Perhaps Microsoft wanted the ability to include any last minute changes without redoing the interface.

The following list defines the actual algorithms used:

- *No Encryption*: users can connect using no encryption at all. When this option is not checked, all connections must be encrypted or users are not allowed to connect. Obviously, leaving this option enabled is important unless you are sure you want to reject connections that use no encryption.
- *Basic*: allows connections using 40-bit connections, including forms of DES for IPSec or Microsoft Point-to-Point Encryption (MPPE) for PPTP.

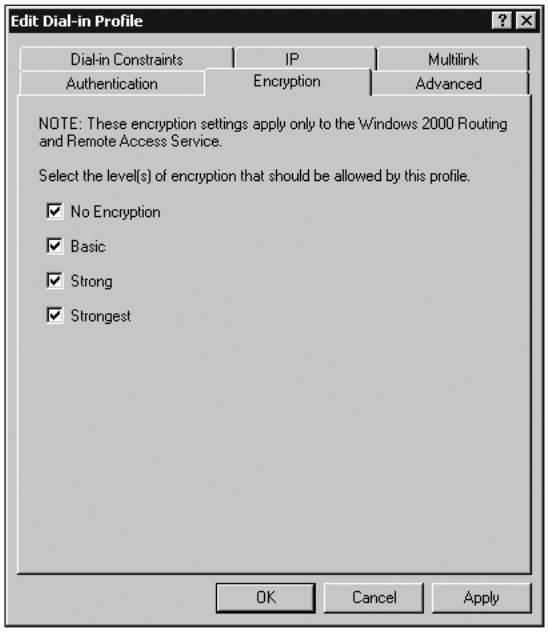

Figure 6-16 Encryption page for a profile

- *Strong*: allows connections using 56-bit encryption, also including forms of DES for IPSec or Microsoft Point-to-Point Encryption (MPPE) for PPTP.
- *Strongest*: allows connections using 128-bit encryption, including triple DES for IPSec and the 128-bit version of MMPE for PPTP.

At the end of the chapter, Hands-on Project 6-5 walks you through the step-by-step process of enabling encryption protocols.

Advanced Properties You use the Advanced page mainly to configure the RRAS server to interact with a RADIUS server. The page lets you add specific attributes (some defined in the RADIUS standard and some for particular vendors) to incorporate into the profile.

For the exam, it's not necessary that you know much detail on the attributes available through the Advanced page of the profile properties. It is enough to know the Advanced page is where you add additional attributes.

Configuring a Virtual Private Networking Connection

Virtual Private Networking (VPN) offers a way to create a logical connection between two computers over an existing IP routing infrastructure. This means that two computers connected by a public network like the Internet can create an additional private connection between them that runs TCP/IP or any other supported protocol and also supports authentication and encryption.

VPNs are typically used in one of two contexts:

- To connect a client to a VPN server. A common scenario is a remote user that first connects to the Internet via a local ISP and then establishes an additional, virtual connection over the Internet to a VPN server on the company network.
- To connect two VPN servers. A common scenario is a company with two locations (and therefore two LANs) that each have Internet access and an RRAS server configured for use with VPN. You can configure these servers to route messages between one another over the Internet using VPN.

In both of these contexts, the main reason you might use VPN instead of traditional dial-up access is simply cost. If you have remote users in a separate calling area from the main network or two networks separated by distance, connecting to the Internet locally instead of through long-distance calls means pretty good savings. Other features also make VPN attractive: it is often easier to configure and more secure than dial-up solutions.

VPN Components

Several components make up a complete VPN solution. These include:

- A VPN Server
- A VPN Client
- A connection between the client and server (VPN connection)
- VPN protocols

VPN Server

For the purposes of our discussion, a VPN server is a Windows 2000 server running the Routing and Remote Access Service configured to support VPN connections. In addition, the server typically has one connection to the Internet and a separate connection to the local network. When you enable RRAS on a server, it is automatically configured to support VPN ports. All you have to do is configure them. You can also specify a server to become a VPN server during the process of enabling RRAS.

VPN Client

A VPN client is any computer that can initiate a VPN connection to a VPN server. This client may be a remote user connecting to a main network or a router connecting to another router. Most operating systems have some sort of VPN client available, even if the operating systems

themselves do not come with built-in support. Windows 98, Windows ME, Windows NT 4.0, and Windows 2000 all include built-in support for use as a VPN client.

VPN Connection

The routing infrastructure for a VPN connection must be some form of IP network, whether this network is the Internet or a private IP network. Often referred to as the **transit internetwork**, this network serves as the basis for the VPN connection. Once the client and server are both connected to the transit internetwork, the client can use TCP/IP or other networking protocols that it shares with the server to establish the VPN connection. For example, if the main company network uses IPX as its primary network protocol, the client probably wants to establish a VPN connection using IPX.

This capability of VPNs to use one networking protocol on top of another is often called tunneling; the virtual network tunnels through the actual network. Using a tunneling protocol supports the ability to tunnel; both sides of the connection use the protocol to create, monitor, and maintain the virtual network. Windows 2000 supports two tunneling protocols: Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP).

Point-to-Point Tunneling Protocol (PPTP) An extension of the PPP remote access protocol, **Point-to-Point Tunneling Protocol** uses a TCP connection for tunnel maintenance and allows IP, IPX, or NetBEUI traffic to be encrypted and then encapsulated within an IP header that can be sent across the IP internetwork.

Layer 2 Tunneling Protocol (L2TP) **Layer 2 Tunneling Protocol** combines PPTP and another protocol called Layer 2 Forwarding. Although the L2TP specifications support transit internetworks using IP, X.25, Frame Relay, or ATM, Windows 2000 only supports L2TP over IP. One distinct advantage that L2TP has over PPTP is that L2TP supports both authentication and encryption for a connection, while PPTP supports only encryption. In addition, L2TP is always used with IPSec, a more secure encryption mode than that used by PPTP, Microsoft Point-to-Point Encryption (MPPE).

Installing and Configuring a VPN Server

To act as a VPN server, a computer must have a permanent and dedicated link to the Internet or to whatever IP network you create the VPN on. Otherwise, the client cannot initiate a connection whenever it needs one. If you already installed RRAS on a server, then that server is already configured to use VPN; you may just not know about it. If you have not yet installed RRAS, you can enable it on your server and specify that it be used as a VPN server. This section discusses both methods.

Installing RRAS as a VPN Server

If you do not yet have RRAS enabled on your server, you need to enable it, activate it, and configure it for use with VPN. This procedure is relatively simple and, for the most part, the same as the procedure for enabling RRAS as an RRAS server. This chapter discussed this

procedure in detail earlier, and there's really not much need to go over it again. The one difference is that when you come to the Common Configurations page of the wizard (refer to Figure 6-2 for a refresher), you should select the Virtual Private Network (VPN) Server option instead of the RRAS Server option. Once you complete the wizard's steps, your new VPN server is ready to accept connections.

Using VPN on an Existing RRAS Server

If you already enabled RRAS on your server and chose something besides the VPN option on the Common Configurations page, you can configure it as a VPN server without having to reinstall the service. All you have to do is open the property pages through the tabs for the server in the RRAS snap-in and make sure that the Remote Access Server option on the General page is enabled. Once you do this, your server can accept VPN connections.

Configuring VPN Ports

Just like enabling VPN on an RRAS server, configuring VPN is a pretty simple process. If you enabled VPN on the server using one of the two methods just discussed, then you have a pretty functional VPN server right off the bat. You can customize a few settings, however.

VPN is primarily managed through the Ports container in the RRAS snap-in. (It's under the server in the left pane.) When you select this container, shown in Figure 6-17, you see a number of objects in the right pane named WAN Miniport. Each represents a virtual port; each port supports either PPTP or L2TP. RRAS is configured by default to accept up to five connections of each type, and these default connections are numbered 0 through 4. Thus, the complete name of a port might be *WAN Miniport (L2TP)(VPN3-1)*.

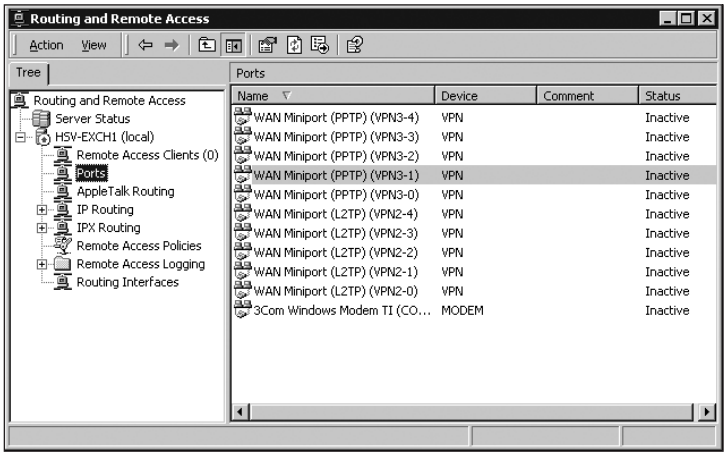

Figure 6-17 Ports container

Columns to the right of the device name list the type of port (in this case VPN) and the status of the port (active or inactive). You can also see a detailed status page for a port by right-clicking the port and choosing Status from the shortcut menu.

To configure settings for the ports on your system, right-click the Ports container itself and choose Properties from the shortcut menu. This opens the Ports Properties dialog box shown in Figure 6-18.

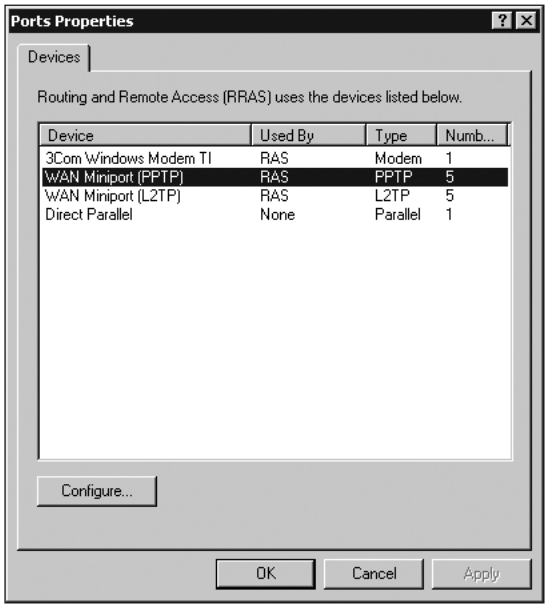

Figure 6-18 Properties for the Ports container

Notice that the dialog box lists both port types (PPTP and L2TP) and indicates the number of ports of each kind. To configure the properties for a port type, select it from the list and click the Configure button. This opens the Configure Device dialog box shown in Figure 6-19. Using this dialog box, you can specify whether the port type may accept incoming connections. Disabling this option essentially turns off the ports of that type. You can also specify whether the port type can be used for demand-dial connections. You need to disable this if you do not want your server to be able to connect to other servers using the port type. You can use the Phone Number field to enter the IP address of the public interface VPN clients use to connect. This would be necessary, for instance, if you had policies in place that granted or denied access based on the number dialed by the client. Finally, you can use this dialog to indicate the number of ports you want available for the port type. By default, you get five of each type, but you can set the number of ports to any number from 0 through 1000.

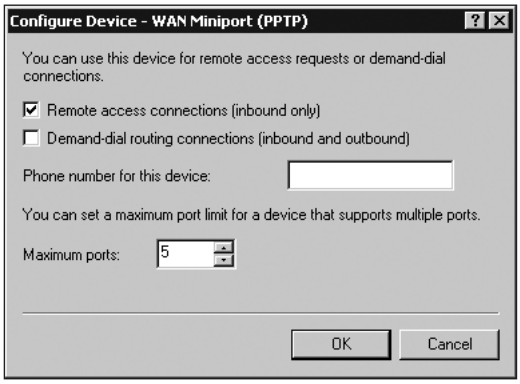

Figure 6-19 Configuring a device from the Ports property pages

Configuring a VPN Demand-dial Interface

A demand-dial interface enables your server to connect to another router or VPN server whenever it needs in order to route information. Creating a demand-dial interface is fairly straightforward, but you need to know a few things before you get started:

- The name and IP address of the router to which you will connect
- The tunneling protocol (PPTP or L2TP) supported by the other router
- A username and password so that the server can connect to the other router

You can practice setting up a demand-dial interface in Hands-on Project 6-7 at the end of the chapter.

Configuring RAS for DHCP Integration

Dynamic Host Configuration Protocol (DHCP) automatically assigns IP addresses and other TCP/IP configuration parameters to clients on a network. This eliminates the administrative hassle of manually configuring every client with an IP address, a subnet mask, a default gateway, DNS and WINS information, and more—a time-consuming and often error-producing process. DHCP is designed to allow clients to broadcast requests for information received by DHCP servers on the same IP network. If clients and servers are on different IP networks, one device on the clients' network must act as a **DHCP Relay Agent** that receives client requests and forwards them to an appropriate DHCP server on another network. This means that a client must be on the same network with either a DHCP Server or a DHCP Relay Agent in order to use DHCP.

When deploying remote access, you must make some decisions about how to handle IP addressing for remote clients. Put simply, you have three choices:

- You can configure your clients with static IP addresses by going to the actual computer. This is usually only a good choice when you are configuring a relatively small number of clients.

- You can configure your RRAS Server as a DHCP Server. If your network is fairly small or if it is easy to deploy an additional DHCP Server, this choice is nice because it lets the RRAS Server assign IP addresses from its own address pool. For more information on configuring a Windows 2000 Server as a DHCP Server, see Chapter 3.
- You can configure your RRAS Server as a DHCP Relay Agent. This is a good option if your RRAS Server is on a separate network from your DHCP Server and your existing DHCP infrastructure makes it hard to simply configure the RRAS Server as a DHCP Server. This section focuses on this last option.

Installing the DHCP Relay Agent

The process of installing a DHCP Relay Agent on your RRAS server is done within the RRAS snap-in. You install it as a new protocol by right-clicking the IP Routing container on the appropriate RRAS server and selecting the New Routing Protocol from the shortcut menu. Hands-on Project 6-8 at the end of the chapter details this process.

Before you can install the DHCP Relay Agent, however, you need to be aware of two things:

- You cannot install a DHCP Relay Agent on a computer that already acts as a DHCP Server.
- You cannot install a DHCP Relay Agent on a computer that runs the Network Address Translation (NAT) protocol.

If you meet these requirements, however, you can install and begin configuring the DHCP Relay Agent.

Configuring the DHCP Relay Agent

Once you install the DHCP Relay Agent, you configure it from two different places. The first is the property pages of the DHCP Relay Agent itself. The second is on the actual interface to which the agent is attached.

Configuring DHCP Relay Agent Properties

Right-click DHCP Relay Agent inside the IP Routing container, and select Properties from the shortcut menu to open the General page for the agent, shown in Figure 6-20. You can configure only one setting for the agent, and that is to what specific DHCP servers the agent points. Just enter an appropriate IP address in the Server address field, and click Add. When you configure all of the DHCP Servers, click OK to finish.

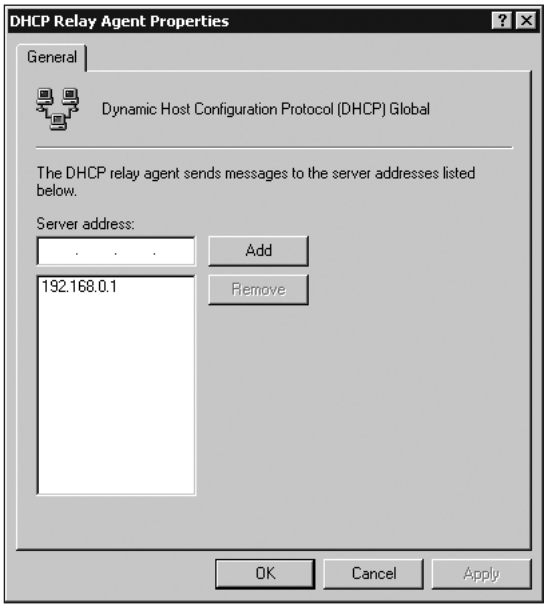

Figure 6-20 Property page for the DHCP Relay Agent

Configuring a Specific Interface

Once you configure the list of DHCP servers to which the agent can forward requests, you must attach the agent to the specific network interfaces that the agent will use. To attach an interface, right-click the DHCP Relay Agent object and select New Interface from the shortcut menu. A dialog box opens that lists all of the interfaces configured on the server. Choose the interface you want, and click OK to make the attachment.

Once you make the attachment, a property page for the interface, shown in Figure 6-21, opens automatically. You can also open these properties later by right-clicking the object for the interface (stored in the Routing Interfaces container for a server) and choosing Properties from the shortcut menu. Make sure the Relay DHCP packets option is enabled if you want the interface used for forwarding DHCP requests. The hop-count number controls the number of additional routers that a request is allowed to pass through, and the Boot threshold number controls how long the interface waits before forwarding any requests it receives.

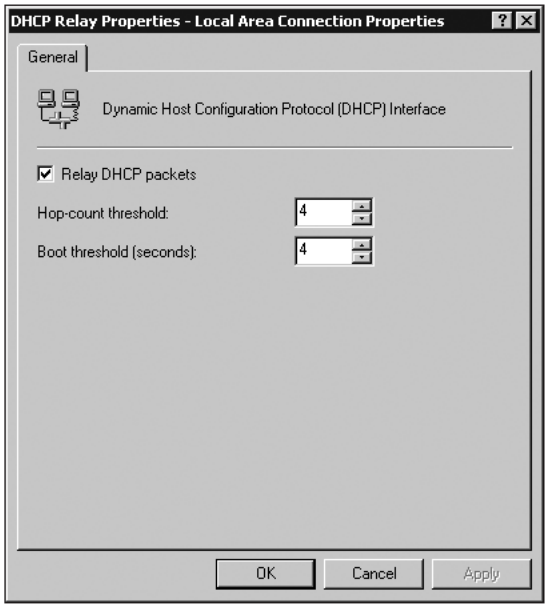

Figure 6-21 General property page for a new interface

Managing, Monitoring, and Troubleshooting RAS

You manage and monitor an RRAS Server using several tools. You can monitor general server and port activity using the RRAS snap-in. You can also use the snap-in to configure logging for the RRAS Server. **Net Shell (netsh)** is a command-line tool used to configure and monitor Windows 2000 networking components, including RRAS. Finally, **Network Monitor** is a powerful application provided with Windows 2000 Server that allows you to capture and examine network packets going in and out of a server for troubleshooting purposes.

Monitoring Server Activity

Just above the list of servers in the RRAS snap-in is an object named Server Status. When you select this item, the details of all servers you are configured to administer show in the right pane. For each server, you receive information on the status of the server (started or stopped), the kind of server, the number of ports configured on it, the number of ports in use, and how long the server has been up. This window provides a snapshot of overall server activity.

Monitoring Ports

In addition to monitoring the status of servers, you can also view the status of each port on a server. To view the status of a port, right-click the port (you'll find it in the Ports container on a server) and select Properties from the shortcut menu. The dialog box that opens shows

the line speed of the port, the amount of data transmitted and received over the port, and the network address for each protocol configured for use on the port. This tool provides a good way to determine whether or not a port is active and how much it is used.

Logging

If you recall, one of the property pages for an RRAS server is named Event Logging. Figure 6-22 shows this page on which you configure the level at which logging occurs for the RRAS server.

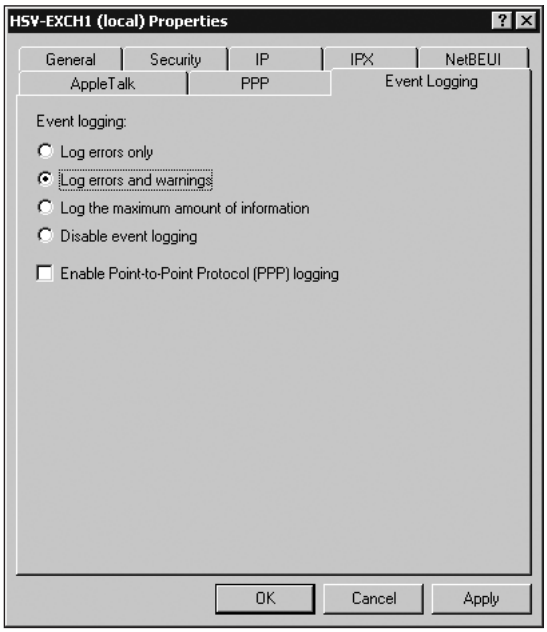

Figure 6-22 Event Logging property page for a server

The options on the Event Logging property page include:

- *Log errors only*: logs only the errors that occur on the RRAS server.
- *Log errors and warnings*: logs errors and warning messages. This is the default choice and usually also the best choice because logging errors and warnings give you good information about problems after they happen and often give fair warning of problems before they happen.
- *Log the maximum amount of information*: logs basically everything that happens on the server. This includes errors, warnings, and even information messages for successful events. This option causes the RRAS service to log a huge amount of messages and is usually only a good choice when you're troubleshooting a particular problem. Turn it on while troubleshooting, and be sure to turn it off when you finish.

- *Disable event logging*: is pretty self-explanatory.
- *Enable Point-to-Point Protocol (PPP) logging*: logs all messages regarding PPP connections to the server. This option is also useful for troubleshooting but can fill an event log pretty quickly if left on for an extended period.

Each RRAS server also has a container under it in the RRAS snap-in named Remote Access Logging. Within this folder are objects representing the actual log files stored for the server. Right-click any log file and choose Properties from the shortcut menu to further configure logging for that server. Figure 6-23 shows the Settings page with these properties. On this page, you can enable or disable the logging of accounting and authentication requests, as well as periodic status for the server.

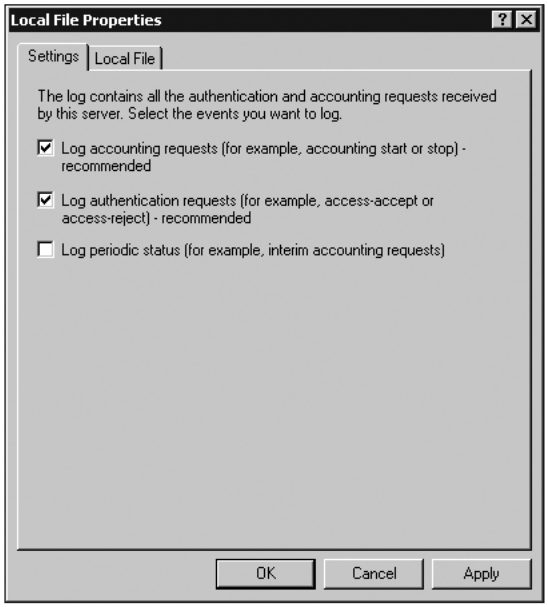

Figure 6-23 Settings property page for a log file

The Local File tab of the Local File property pages, shown in Figure 6-24, controls the physical aspects of how the file is written to disk. You can change the format of the log between the commonly used Internet Authentication Service (IAS) Format and a database compatible format, which makes it easy to import the log file into a database for processing. You can also change the period at which new log files are created and the directory for storing the file.

Using the Net Shell Tool

The Net Shell (usually called netsh) tool is a command-line and scripting tool that lets you configure and monitor Windows 2000 networking components. It is automatically installed with Windows 2000 in the \system32 subfolder of the Windows folder.

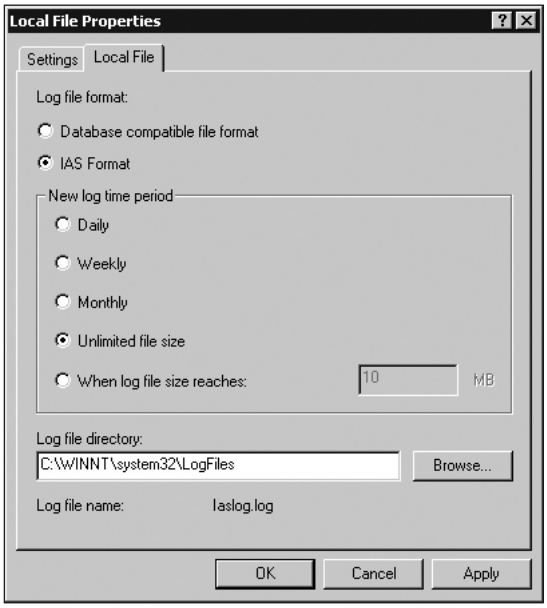

Figure 6-24 Local File property page for a log file

When you run netsh, you enter a shell that accepts special netsh commands. You can run this shell in one of two modes. In online mode, commands execute as soon as you type them into the shell. Offline mode saves commands as you type them and then executes them in batches when you use a special commit command.

With regard to RRAS, netsh provides the ability to access certain RRAS configuration settings, routing settings, and interface settings. You can learn more about the specifics of using the tool by typing netsh /? after the command prompt.

Using Network Monitor

Network Monitor is an application that comes with Windows 2000 and allows you to capture and view the actual packets of information being transmitted over a network interface. Once you capture a set period of traffic, you can filter your view of that traffic so that you can examine packets from, say, a specific protocol or time period. Using Network Monitor enables you to build a solid picture of network traffic patterns and to spot potential problems before they occur. You learn more about using Network Monitor in Chapter 7.

Chapter Summary

- The RRAS Service runs on a Windows 2000 server and enables other servers or client computers that are not connected to the network via a permanent cable to establish temporary connections over phone lines, ISDN lines, or services such as X.25. Once a

computer establishes a connection with the RRAS server, that computer can access the resources on the RRAS server and possibly access the other computers on the same network as the server, depending on the server's configuration.

- RRAS provides users with two types of remote access connections: dial-up networking connections, in which the user dials in to the RRAS Server using a phone or ISDN line, and Virtual Private Networking (VPN) connections, in which the user connects first to a transit internetwork such as the Internet and then establishes a virtual connection over that internetwork to the RRAS server.
- Two types of protocols govern the transmission of information between a client and an RRAS server. Remote access protocols control how a dial-up connection is actually established. (In the case of a VPN connection, tunneling protocols like PPTP and L2TP do the same thing.) The PPP protocol is the most common protocol in use today. Networking protocols determine how data is segmented and shaped for transmission over the connection once it is established. RRAS supports the IP, IPX, NetBEUI, and AppleTalk networking protocols.
- RRAS is installed by default along with Windows 2000 but not enabled. To enable it, you must use the Routing and Remote Access Server snap-in located on the Start menu. Once enabled, most RRAS configuration also happens within this snap-in. You use it to configure server properties, monitor servers and ports, create interfaces, and install new protocols. You also use it to create new remote access policies and profiles that determine what users may and may not access the system.
- The other primary tool for configuring RRAS is the Active Directory Users and Computers snap-in used for creating and managing user profiles. For each individual user, a Dial-in property page lets you configure whether the user can use remote access and several other settings that govern that access.
- Remote access security comes primarily in the form of authentication of a user's credentials. RRAS supports a number of authentication methods, including Password Authentication Protocol (PAP), Shiva Password Authentication Protocol (SPAP), Challenge Handshake Authentication Protocol (CHAP), Microsoft CHAP (MS-CHAP), and Extensible Authentication Protocol (EAP). In addition to authentication, RRAS also provides ways to control the connections made to the system and to control the access of individual users.
- Several tools are used in managing and monitoring an RRAS server. You can monitor general server and port activity using the RRAS snap-in. You can also use the snap-in to configure logging for the RRAS Server. Net Shell (netsh) is a command-line tool used to configure and monitor Windows 2000 networking components, including RRAS. Finally, Network Monitor is a powerful application provided with Windows 2000 Server that allows you to capture and examine network packets going in and out of a server for troubleshooting purposes.

KEY TERMS

Accounting provider — Server (typically a RADIUS server) that logs the activity and connection time for a remote user. This is often used to charge remote clients for online time, as in the case of an ISP providing Internet service.

Active Directory Users and Computers — Tool used to configure the objects in the Windows 2000 Active Directory. Among other things, you use this tool to configure the properties of user accounts. Dial-in properties for a user include whether the user may dial in to the RRAS server and whether a callback number should be used.

authentication — Process of verifying a user's credentials so that the user may log on to the system. Authentication is normally performed using a username and password. Authentication may be unencrypted (clear text) or use any of a number of **encryption** types.

Bandwidth Allocation Control Protocol (BACP) — *See* Bandwidth Allocation Protocol (BAP).

Bandwidth Allocation Protocol (BAP) — Together with the Bandwidth Allocation Control Protocol (BACP), allows a client to add and remove links dynamically during a multilink session to adjust for changes in bandwidth needs.

Challenge Handshake Authentication Protocol (CHAP) — Type of authentication in which the authentication agent sends the client program a key for encrypting the username and password.

demand-dial routing — Allows an RRAS server configured as a router to dial-up a remote router whenever it needs to send messages to that router.

DHCP Relay Agent — When they boot up, DHCP clients broadcast a message to their local IP subnet looking for a DHCP server to provide them with IP addressing information. These broadcast messages typically do not pass through routers. One way to avoid putting a full DHCP server on every subnet is to configure a computer as a DHCP Relay Agent. This computer intercepts the DHCP client requests and forwards them across the router to the DHCP server. RRAS has the capability to serve as a DHCP Relay Agent.

Dial-Up Networking — Name given to the process and interface that most versions of Microsoft Windows use to dial in to a remote server.

Dynamic Host Configuration Protocol (DHCP) — Protocol used to automatically assign IP addressing and other TCP/IP information to clients. DHCP is considered easier and more reliable than manual addressing.

Encryption — Process of translating information into an unreadable code that can only be translated back (decrypted) by using a secret key or password.

event logging — Most applications in Windows (and Windows itself) log events to a file. Events are bits of information and any errors generated by these applications. Once logged, you can view the events using the Event Viewer utility.

Extensible Authentication Protocol (EAP) — General protocol for PPP authentication that supports multiple authentication mechanisms. Instead of selecting a single authentication method for a connection, EAP can negotiate an authentication method at connect time.

Internet Group Management Protocol (IGMP) — Standard protocol for IP multicasting over the Internet. It is used to establish host memberships in particular multicast groups.

IP Security (IPSec) — Set of protocols that supports the secure exchange of data at the IP layer. In RRAS, IPSec is used in conjunction with L2TP in the formation of Virtual Private Networks.

IPX (Internetwork Packet eXchange) — Networking protocol developed by Novell for use primarily with their NetWare operating systems. Since NetWare is such a popular network operating system, most other operating systems, such as Microsoft Windows, provide an IPX-compatible networking protocol. In Windows 2000, this IPX-compatible protocol is named NWLink.

Layer-Two Tunneling Protocol (L2TP) — Extension of the PPP remote access protocol; one type of tunneling protocol used to form Virtual Private Networks.

Link control protocol (LCP) — LCP extensions include a number of enhancements to the LCP protocol used to establish a PPP link and control its settings. One of the primary enhancements included is the ability for the client and server to agree dynamically on protocols used on the connection.

Microsoft CHAP (MS-CHAP) — Modified version of CHAP that allows the use of Windows 2000 authentication information. There are two versions of MS-CHAP. Version 2 is the most secure, and all Microsoft operating systems support it. Other operating systems sometimes support version 1.

multicast routing — Targeted form of broadcasting that sends messages to a select group of users instead of all users on a subnet.

Multilink Protocol (MP) — Used to combine multiple physical links into a single logical link. For example, you could use MP to combine two 56-KB modem links into a 128-KB link.

NetBEUI — Enhanced version of the NetBIOS networking protocol primarily used on older versions of Microsoft and IBM operating systems.

Net Shell (netsh) — Command-line tool used to configure and monitor Windows 2000 networking components, including RRAS.

Network Address Translation (NAT) — Router standard that translates IP addresses on a private network into valid Internet IP addresses. NAT makes it possible for a single computer with Internet connectivity to share its Internet connection with other computers on the network through a single IP address.

Network Monitor — Tool that comes with Windows 2000 and allows you to capture and view data packets passing over the network.

networking protocols — Standard language used by two computers to communicate over a network. Networking protocols define how information is fragmented and shaped for passage over the network.

Password Authentication Protocol (PAP) — Authentication method that transmits a user's name and password over a network and compares them to a table of name-password pairs.

Point-to-Point Protocol (PPP) — Remote-access protocol used to establish a connection between two remote computers. RRAS supports PPP for dialing both in and out.

remote access — Broadly defines the ability of one computer to connect to another computer over a dial-up or other WAN connection and to access resources remotely.

remote access policy — Used to configure conditions under which users may connect using a specific remote access connection. You can include restrictions based on criteria such as time of day, type of connection, authentication, and even length of connection.

remote access profile — Associated with policies and containing settings that determine what happens during call set up and completion.

remote access protocols — Define the way in which one computer connects to another computer over a WAN link. PPP and SLIP are the two main remote access protocols in use today, though the newer and stronger PPP is much more common.

Remote Authentication Dial-In User Support (RADIUS) — Authentication and accounting system used by many ISPs to verify user credentials and log user activity while the user is connected to a remote system.

remote control — Process in which a client computer connects to a remote server and actually takes control over that server in a separate window on the client computer. Activities within this window seem to occur as if the user is actually sitting at the server computer. All applications run on the server. RRAS does not support remote control, only remote access.

router — Device used to connect different IP subnets and to route data between them.

Routing and Remote Access Service (RRAS) — Windows 2000 service that provides remote access and routing functionality to remote clients.

Serial Line Interface Protocol (SLIP) — Older protocol developed in UNIX and still in wide use today. Windows 2000 RRAS supports SLIP in dial-out configurations, but you cannot use a SLIP client to dial in to an RRAS server.

Shiva Password Authentication Protocol (SPAP) — Included mainly for compatibility with remote access hardware devices manufactured by Shiva, a private company now owned by Intel. SPAP isn't really used much on most networks.

transit internetwork — Basic IP infrastructure over which a Virtual Private Network is created. Typically, the transit internetwork is the Internet itself, though other IP networks may be the transit internetwork.

Transmission Control Protocol/Internet Protocol (TCP/IP) — Suite of networking protocols designed to transfer data between computers on the Internet. TCP/IP is becoming the most popular networking protocol used on private networks, as well.

user profile — Information associated with a user account. Profiles of users who are members of a Windows 2000 domain are stored in the Active Directory, and profiles of users who are not members of a domain are stored on the local computer.

Virtual Private Networking (VPN) — Secure, logical network constructed directly between a VPN client and a VPN server on top of a physical transit internetwork such as the Internet.

Review Questions

1. Which of the following must be true in order to use remote access policies?
   a. The RRAS server must not be a domain controller.
   b. The RRAS server must be configured as a DHCP Server or DHCP Relay Agent.
   c. Active Directory must be running in native mode.
   d. RRAS must be running on Windows 2000 Advanced Server or Datacenter Server.
2. Which of the following are valid tunneling protocols for Virtual Private Networks?
   a. PPP
   b. PPTP
   c. L2TP
   d. IPSec
3. You need to configure callbacks for a group of remote users. Though you could do this by enabling callbacks for every individual user's profile, what is a better solution?
   a. Create a Remote Access Policy for the group.
   b. Create a remote access profile for the group.
   c. Enable callbacks for all users in the RRAS snap-in.
   d. Use DHCP to assign callback information.
4. The _____ protocol allows the sharing of a single Internet connection with other computers on a LAN.
5. The Windows 2000 implementation of L2TP supports access over IP, X.25, and ATM. True or false?
6. Which of the following remote access protocols can Windows 2000 use to accept incoming calls?
   a. PPP
   b. PPTP
   c. SLIP
   d. IP
7. An RRAS server can only support one networking protocol on a system. True or false?
8. You want your clients to be assigned IP addressing information automatically. Your network uses a DHCP system, but you configured your RRAS Server on a different IP network than the rest of your network. Which of the following solutions lets the existing DHCP servers assign addresses to the remote clients?
   a. Install a new demand-dial interface on the RRAS server.
   b. Configure the RRAS server as a DHCP server, and set it to forward requests.
   c. Install the DHCP Relay Agent protocol on the RRAS server.
   d. Install a WINS proxy agent on the RRAS server.

9. Which of the following statements is true?
   a. Settings made on a user's profile override settings made in a Remote Access Policy applied to that user.
   b. Settings made in a Remote Access Policy applied to a user override settings made on that user's profile.
   c. The Remote Access Profile determines conflicting settings in a user's profile and any remote access policy applied to that user.
   d. You can use a server's properties in the RRAS snap-in to specify whether user profiles override remote access policies or vice versa.
10. Conventional dial-up connections are much easier to configure and manage than VPN connections. True or false?
11. You can use the _____ tool to capture and view packets of information transferred over the network.
12. In the RRAS snap-in, what is the allowed number of incoming VPN connections configured through?
    a. Ports container
    b. Connections container
    c. Interfaces container
    d. VPN container
13. You want to ensure that all user authentication information passed between remote clients and an RRAS server is encrypted. Which of the following authentication methods should you disable?
    a. PAP
    b. CHAP
    c. SPAP
    d. EAP MD5-CHAP
14. What type of encryption does PPTP use?
    a. IPSec
    b. PAP
    c. MMPE
    d. DES
15. _____ is a targeted form of network broadcasting that sends information to a select group of users instead of all users connected to a network.

16. Which of the following is true of the DHCP Relay Agent protocol? (Choose all that apply.)
    a. It is used only on RRAS servers that use static addressing for clients.
    b. All RRAS servers require it.
    c. Servers running the DHCP service cannot use it.
    d. Servers running NAT cannot use it.
17. You want to assign a user the same IP address every time the user dials in to the RRAS server. How can you do this?
    a. Set up the DHCP Relay Agent on the server, and configure a reserved address for the user.
    b. Enter the IP address using the static IP option on the user's property pages in Active Directory Users and Computers.
    c. Configure a reserved address for the user with the property pages of the IP Routing container in the RRAS snap-in.
    d. You cannot do this.
18. You can create a remote access profile without association to a Remote Access Policy. True or false?
19. Your RRAS server currently lets log files grow to unlimited size. You want to configure it so that it creates new log files every week. Where can you do this?
    a. The Settings tab of the server's property pages
    b. The Event Logging tab of the server's property pages
    c. The Local File tab of the log file's property pages
    d. The IP tab of the log file's property pages
20. Stopping the RAS service causes the service to erase its settings. True or false?

HANDS-ON PROJECTS

All Hands-on Projects in this chapter require at least one server computer set-up as described in the lab setup section in the front of this book. To complete these exercises, you must have completed the projects in Chapters 2 and 3 on installing networking protocols and configuring DHCP.

Project 6-1

To install Routing and Remote Access Service:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.
2. Right-click the computer name and select **Configure and Enable Routing and Remote Access** from the context menu.

3. On the Welcome screen for the Routing and Remote Access Server Setup Wizard, click **Next**.
4. On the Common Configurations screen, select the **Remote Access Server** option and then click **Next**.
5. On the Remote Client Protocols screen, under Protocols, make sure that TCP/IP is listed.
6. Verify that the **Yes**, **all the required protocols are on this list** option is selected, and click **Next**.
7. On the IP Address Assignment screen, make sure that the **From a specified range of addresses** option is selected, and then click **Next**.
8. In the Address Range Assignment window, click **New**.
9. In the Starting Address field, type **192.168.0.200**, and in the End Of IP Address field, type **192.168.0.225**.
10. Under Number of Addresses, verify that **26** is the number, click **OK** to close the Edit Address Range window, and then click **Next**.
11. On the Managing Multiple Remote Access Servers screen, verify that the **No**, **I don't want to set this server up to use RADIUS now** option is selected and then click **Next**.
12. Click **Finish**.
13. Click **OK** to respond to any warning messages that appear.

Project 6-2

To configure the multilink protocol for incoming connections:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.
2. Right-click the server for which you want to configure multilink, and select the **Properties** command from the shortcut menu.
3. Click the **PPP** tab.
4. Remove the check mark next to the **Multilink connections** option.
5. Click **OK** to close the property pages.

Project 6-3

To configure the IP protocol to be the only protocol clients can use for remote access:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.
2. Right-click the server you want to configure, and select **Properties** from the shortcut menu.

3. Click the **IP** tab.
4. Make sure that the **Enable IP Routing** and **Allow IP-based remote access and demand-dial connections** options are selected.
5. Click the **IPX** tab, if there is one.
6. Disable the **Allow IPX-based remote access and demand-dial connections** option.
7. Click the **NetBEUI** tab, if there is one.
8. Disable the **Allow NetBEUI-based remote access clients to access** option.
9. Click the **AppleTalk** tab, if there is one.
10. Disable the **Enable AppleTalk remote access** option.
11. Click **OK** to close the server's property pages.

Project 6-4

To create a new remote access policy:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.
2. Right-click **Remote Access Policies** and select **New Remote Access Policy**.
3. Type **All users time restrictions** in the Policy friendly name field, and then click **Next**.
4. Click **Add** to add a condition.
5. Select **Day-and-Time Restrictions** and then click **Add**.
6. On the Time of day constraints dialog box that opens, select the times and days you want to allow or deny access to and then click OK.
7. Click **Add** to add another condition.
8. Select the **Windows Groups entry** and click **Add**.
9. On the Windows Groups dialog box that opens, click **Add**.
10. Select **Domain Users** and then click **Add**.
11. Click **OK** twice to close the window and click the **Groups** dialog box, then click **Next**.
12. Select the **Grant Remote Access** permission, and click **Next**.
13. Click **Finish**.

Project 6-5

To configure encryption protocols for a remote access server:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.
2. In the left-hand pane, expand the server you want to configure and select the **Remote Access Policies** container.

3. Right-click the **Allow access if dial-in permission is enabled** policy, and choose **Properties** from the shortcut menu.
4. Click the **Edit Profile** button.
5. Click the **Encryption** tab.
6. Remove the check mark next to the **No Encryption** option.
7. Make sure that the other encryption policies are all enabled.
8. Click **OK** to exit the dialog box, and click **OK** again to return to the RRAS snap-in main window.

Project 6-6

To configure a Windows 2000 Professional client to access a remote server:

1. Click **Start**, point to **Settings**, and then select **Network and Dial-Up Connections**.
2. Double-click the **Make New Connection** icon.
3. On the Welcome page of the Network Connection Wizard, click **Next**.
4. On the Network Connection Type page, select **Dial-up to private network** and click **Next**.
5. On the next page, type the **phone number** to which you want to connect and if you want to use any established dialing rules, select that option and click **Next**. Then you need to decide whether to share this connection with other users who can log on to the computer. The default is to share the connection.
6. Click **Next**.
7. Type **a name** for the new dial-up connection, and click **Finish** to exit the wizard. The Connect dialog for the new connection opens automatically when you exit the wizard.
8. Enter **your username** and **password**, and then click **Dial** to connect to the remote network. Once connected, you should be prompted to log on to the remote network using your normal network credentials.

Project 6-7

To create a new VPN demand-dial interface:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.
2. Right-click **Routing Interfaces**, select **New Demand-Dial Interface** from the shortcut menu, and click **Next**.
3. In the Interface Name field, type **a name** for the remote router to which you will connect and click **Next**.
4. On the VPN Type screen, select the **Automatic selection** option and then click **Next**.
5. Enter the **IP address** of the router to which you will connect, and then click **Next**.

6. On the Protocols and Security screen, select the **Route IP Packet On This Interface** option and then click **Next**.
7. Enter the **local router name** in the Dial-Out Credentials dialog box.

   This is the username the router will use when connecting to the remote router. This username will match the name of the demand-dial interface on the remote router.
8. Leave the Domain and Password fields blank, and click **Next**.
9. Click **Finish**.

Project 6-8

To configure a DHCP relay agent to work over Routing and Remote Access:

1. Click **Start**, point to **Programs**, point to **Administrative Tools**, and then select **Routing and Remote Access**.
2. Expand the container for the server on which you want to install the DHCP Relay Agent.
3. Right-click the **General** container under IP routing, and select the **New Routing Protocol** option from the shortcut menu.
4. Choose **DHCP Relay Agent** and then click **OK**.
5. In the IP Routing container, select the **DHCP Relay Agent** object and right-click **Properties** on the shortcut menu.
6. Use the dialog box that opens to configure the IP addresses of DHCP servers.
7. Click **OK** to close the DHCP Relay Agent Properties dialog box.

CASE PROJECTS

Case 1

Figure 6-25 shows the configuration of your company network. Your RRAS server sits on a separate IP subnet outside the perimeter of the rest of the company network. Until now, you had your RRAS server configured to assign IP addresses to remote clients from a pool of addresses you allocated to the server. Recently, many more users are accessing the network remotely and managing IP addresses is becoming harder. Given the layout of the network presented in Figure 6-25, what are your options?

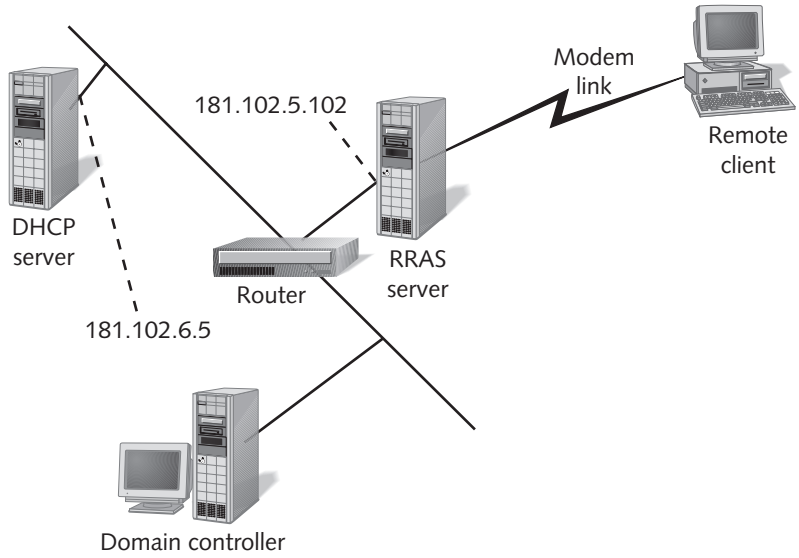

Figure 6-25 Case 1 network diagram

Case 2

The number of remote users in your company continues to grow and the number of users who connect to the company while traveling is also growing quite large. Quite frankly, your company does not want to pay the long-distance fees involved in allowing employees to connect from all over the country for extended periods of time. You mentioned to your manager that implementing a VPN solution using RRAS might be just the thing and she wants to know all the details. Sketch a diagram (using the diagram in Figure 6-25 as your basis, if you like) that shows the following details:

- The RRAS connection to the Internet
- The client at the other end of the Internet connection
- All protocols to be used, including the IP protocol for the Internet, the remote access protocol used by the client, the tunneling protocol used for the VPN connection, and the encryption protocol used by the chosen tunneling protocol
- How the VPN works and why it will be secure

Case 3

You want to create a series of policies that accept or reject users' dial-in attempts based on the following criteria:

- You use VPN as your only remote access method and want to accept only L2TP connections.
- Between 6:00 a.m. and 5:00 p.m., you want all users who have dial-in permissions to be able to dial in.

- Between 5:00 p.m. and 1:00 a.m., you want all users who are members of the Domain Admins, Executives, Engineers, and Marketing security groups to be able to dial in, but no one else.
- Between 1:00 a.m. and 6:00 a.m., you want only members of the Domain Admins group to be able to dial in.
- Using the conditions listed in Table 6-1 earlier in this chapter, outline the policies you need to create to meet these criteria and their proper order.